Accelerated Strategic Computing Initiative

Program Plan

September 1996

U.S. Department of Energy Defense Programs
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Sandia National Laboratories

Accelerated Strategic Computing Initiative

Program Plan

September 1996

U.S. Department of Energy Defense Programs
Lawrence Livermore National Laboratory
Los Alamos National Laboratory
Sandia National Laboratories

Table of Contents

Executive Summary

Ushering in a New Era

On August 11, 1995, President Clinton announced the United States' intention to pursue a "zero yield" Comprehensive Test Ban Treaty and thus reduce the nuclear danger. This decision ushered in a new era in the way the United States ensures confidence in the safety, performance, and reliability of its nuclear stockpile. The President said in his announcement,

> we can meet the challenge of maintaining our nuclear deterrent under a [comprehensive test ban] through a science-based stockpile stewardship program without nuclear testing.

Furthermore, the President also reaffirmed the United States' decision to halt new nuclear weapon designs. This decision means that the U.S. nuclear weapon stockpile will need to be maintained far beyond its design lifetime.

The Accelerated Strategic Computing Initiative (ASCI) is a critical element needed to shift from test-based confidence to science-based confidence. Specifically, ASCI will accelerate the development of simulation capabilities needed to ensure confidence in the nuclear stockpile—far exceeding what might have been achieved in the absence of a focused initiative.

The ASCI Vision…

Shift promptly from nuclear test-based methods to computation-based methods.

ASCI will create the leading-edge computational modeling and simulation capabilities that are essential for maintaining the safety, reliability, and performance of the U.S. nuclear stockpile and reducing the nuclear danger.

Realizing the Vision

To realize its vision, ASCI will create virtual testing and prototyping capabilities based on advanced weapon codes and high-performance computing. Virtual testing is the use of predictive simulations, based on experimental data, to assess and certify the safety, performance, and reliability of nuclear systems. Today, virtual testing and prototyping exist in rudimentary forms. Dramatic advances in computer technology have made virtual testing and prototyping viable alternatives to traditional nuclear and nonnuclear test-based methods.

ASCI will provide computational and simulation capabilities that will help scientists understand aging weapons, predict when

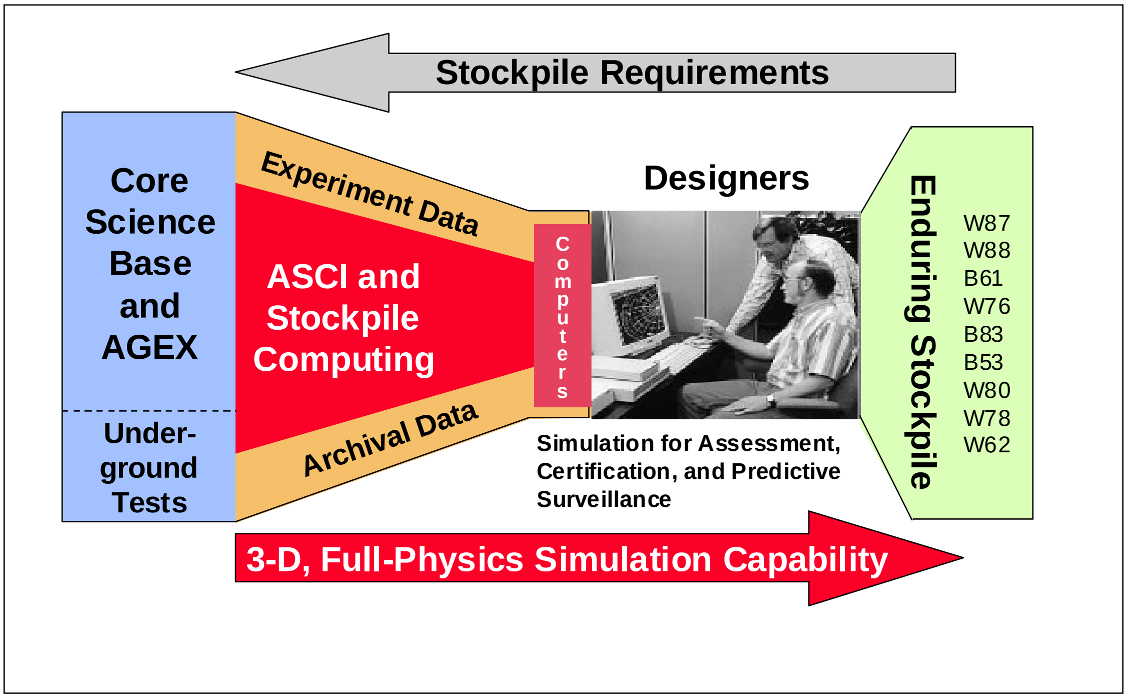

Figure 1. *Simulation, via ASCI, is the only practical, credible path from the known to the unknown for designers to provide an understanding of the enduring stockpile.*

components will have to be replaced, and evaluate the implications of changes in materials and fabrication processes to the design life of the aging weapon systems. This science-based understanding is essential to ensure that changes brought about through aging or remanufacturing will not adversely affect the enduring stockpile.

To meet the needs of stockpile stewardship in the year 2010, ASCI must solve progressively more difficult problems as we move away from nuclear testing. To do this, applications must achieve higher resolution, higher fidelity, three-dimensional, full-physics, and full-system modeling capabilities to reduce empiricism. This level of simulation requires high-performance computing far beyond our current level of performance. A powerful problem-solving environment must also be established to support application development and enable efficient and productive use of the new computing systems.

The ASCI program recognizes that the creation of simulation capabilities needed for virtual testing and prototyping is a significant challenge. This challenge is on par with many aspects of the original Manhattan Project and requires the science and technology resources available only at the national laboratories. This challenge will require close cooperation with the computer industry to accelerate their business plans to provide the computational platforms needed to support ASCI applications. Universities will also play a critical role in advancing the research and development needed for this unprecedented level of simulation.

Enabling Science-Based Stockpile Stewardship

ASCI is a critical element of the Department of Energy's response to the decision ending nuclear testing by enabling the integration of

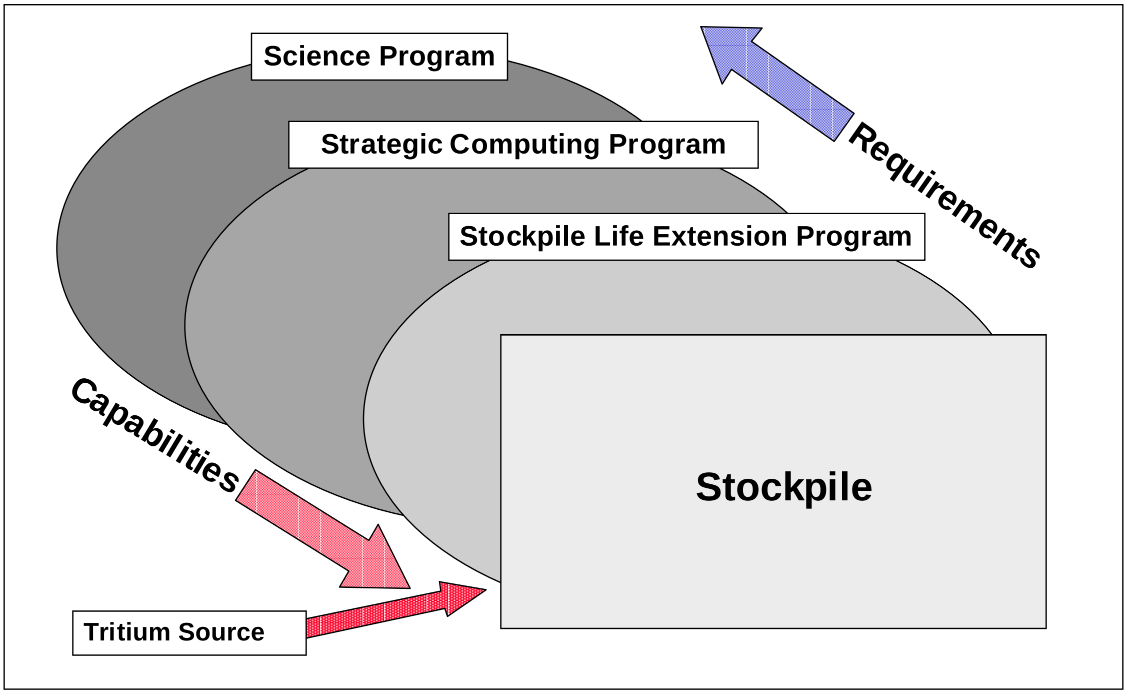

Figure 2. *Support for the enduring stockpile is built on a strong foundation of simulation capabilities, nonnuclear and subcritical test capabilities, and science.*

science into the actual weapons in the stockpile. The Science-Based Stockpile Stewardship (SBSS) program will build on existing means and develop new means to assess the performance of nuclear stockpile systems, predict their safety and reliability, and certify their functionality. The SBSS program not only must respond to the loss of nuclear testing, but also must deal with constraints on nonnuclear testing, the downsizing of production capability, and the cessation of new weapon designs to replace existing weapons. Further complicating matters, weapon components will exceed their design lifetimes, and manufacturing issues and environmental concerns will force changes in fabrication processes and materials of weapon components.

SBSS will support programs responsible for developing the fundamental scientific understanding of nuclear weapons and programs responsible for the surveillance, maintenance, assessment, and certification of the weapons. In the past, much of the integration of the fundamental science development into nuclear weapons was accomplished through testing (and, specifically, underground nuclear tests). In the future, the simulation capabilities provided by ASCI will provide that integration.

The ASCI simulation capabilities will link the experimental data from the aboveground test facilities (AGEX), archival nuclear test data, and improved scientific understanding to provide high-confidence predictive simulation capabilities needed to support decisions about the enduring stockpile. ASCI supports another element of SBSS, the Stockpile Life Extension Program (SLEP), by providing simulation capabilities needed to predict requirements for replacement of aged components and to ensure that those replacements do not introduce new problems into the stockpile. Finally, ASCI will complement and accelerate the ongoing efforts of the Defense Programs core research program for advances

The ASCI Objectives

ASCI has specific program objectives in the areas of performance, safety, reliability, and renewal.

- ***Performance:*** Create predictive simulations of nuclear weapon systems to analyze behavior and assess performance in an environment without nuclear testing.
- ***Safety:*** Predict with high certainty the behavior of full weapon systems in complex accident scenarios.
- ***Reliability:*** Achieve sufficient, validated predictive simulations to extend the lifetime of the stockpile, predict failure mechanisms, and reduce routine maintenance.
- ***Renewal:*** Use virtual prototyping and modeling to understand how new production processes and materials affect performance, safety, reliability, and aging issues. This understanding will help define the right configuration of production and testing facilities necessary for managing the stockpile throughout the next several decades.

These objectives will be realized through the implementation of the five ASCI strategies.

in physics, material sciences, and computational modeling.

The ASCI Strategies

To meet its objectives, the ASCI program has five interrelated strategies.

1. Create Seamless Management: One Program—Three Laboratories

The problems that ASCI will solve for Science-Based Stockpile Stewardship span the activities and responsibilities of the three Defense Programs laboratories (Los Alamos, Sandia, and Lawrence Livermore). Cooperation among the Defense Programs laboratories is essential to solving these problems in an efficient and effective manner. In accordance with this cooperative philosophy, representatives of the laboratories participated in the development of this plan. There has been, and will continue to be, unprecedented cooperation among the three Defense Programs laboratories. The ASCI program will be implemented by project leaders at each of the laboratories, guided by the Office of Strategic Computing and Simulation under the Assistant Secretary for Defense Programs. The weapon laboratories will share ASCI code development, computing, storage, and communication resources across laboratory boundaries in joint development efforts.

2. Focus on Advanced Applications Development

ASCI is developing on an accelerated schedule the progressively higher performance software applications needed to implement virtual testing and prototyping. The key to reaching the Science-Based Stockpile Stewardship objectives outlined for 2010 is our ability to achieve in the intervening years ASCI's critical simulation and applications code milestones. ASCI will provide simulations embodying all the physics needed to predict the safety, reliability, performance, and manufacturability of weapon systems.

It is a formidable challenge to replace the empirical factors and adjustable parameters used in current calculations with predictive physical models. Solving this challenge will require large, complex computer applications codes that drive the scale of computing machinery and the infrastructure. However, increased capability in machinery and infrastructure alone is insufficient. Much of the increased computational capability to be provided by ASCI must come from advances in the applications codes themselves. These applications will integrate 3-D capability, finer spatial resolution, and more accurate and robust physics. Tightly integrated code teams—large interdisciplinary work groups whose objective is to produce coherent software packages for efficient predictive simulations—will develop these codes.

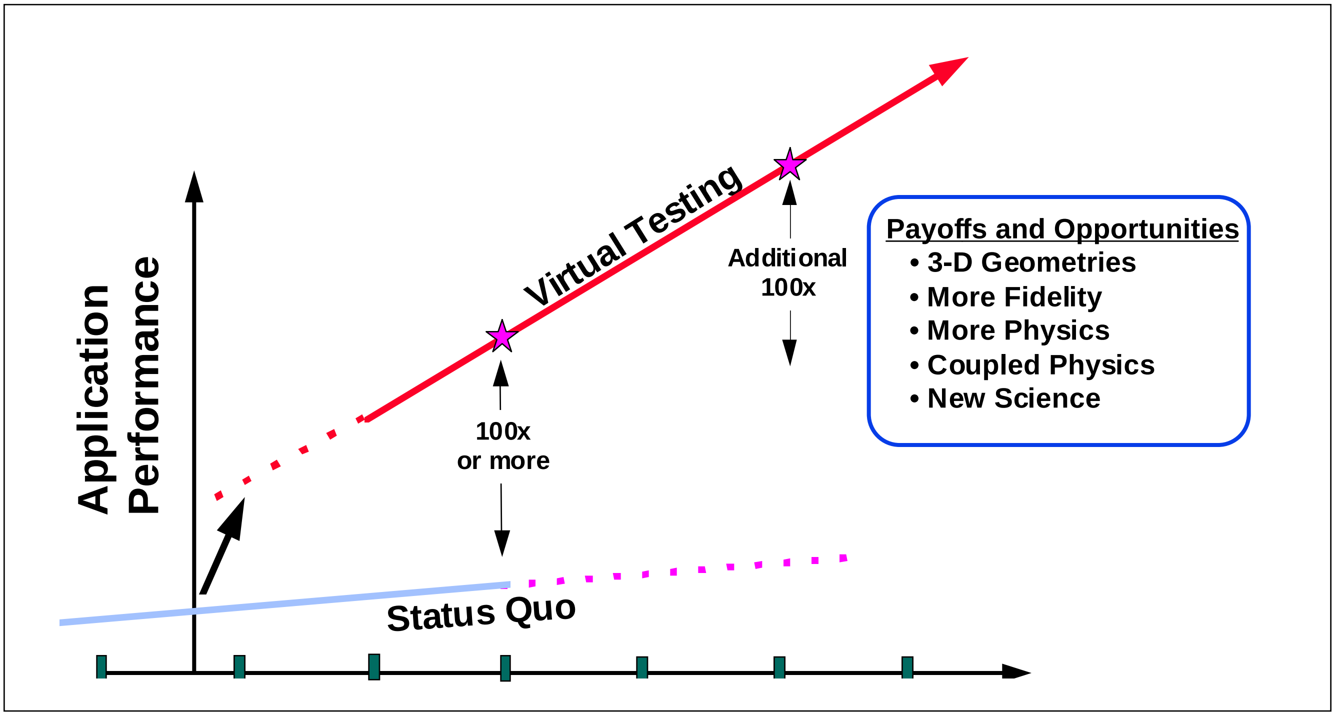

Figure 3. *ASCI will require a significant leap in computational capabilities for virtual testing.*

3. Focus on the High End of Computing

More powerful computers are needed for virtual testing and prototyping applications. ASCI will stimulate the U.S. computing industry to develop high-performance computers with speeds and memory capacities thousands of times greater than currently available models and ten to several hundred times greater than the largest computers likely to result from current development trends. ASCI will partner with various U.S. computer manufacturers to accelerate the development of larger, faster computer systems and software that are required to run Defense Programs applications.

4. Create Problem-Solving Environments

ASCI's unprecedented code development effort will require robust problem-solving computing environments where codes may be developed rapidly. ASCI will develop a computational infrastructure to allow applications to execute efficiently on the ASCI computer platforms and allow accessibility from the weapon designer's desktops. This computational infrastructure will consist of local area networks, wide-area networks, advanced storage facilities, and software development and data visualization tools.

5. Encourage Strategic Alliances and Collaborations

ASCI will require the technical skills of the best scientists and engineers working in academia, industry, and other government agencies in addition to those working in the national laboratories. The need to develop an unprecedented level of simulation capability requires strategic alliances with leading research organizations. These alliances will support the development and credible demonstration of this simulation capability. ASCI will also work with the larger computing community to develop and apply commercially acceptable standards. Finally, ASCI plans to initiate exchange programs to bring top researchers directly into the project while allowing laboratory personnel to expand their experience base in external projects.

Initial Results

This program plan is a revision of one published for comment in April 1995. This revision addresses many of the comments and experiences from the last year. While the ASCI program is carefully planned, it is designed to be flexible in order to respond to changing requirements. ASCI will continue to solicit and address concerns throughout the life of the initiative. Moreover, it has responded to the needs set forth last year—specifically, a “zero yield” Comprehensive Test Ban Treaty and stockpile life extension. Both significant challenges are critical stages that today rely on testing or empirical models and an entirely new activity—the study of the aging process of weapons.

ASCI is a funded program with projects already under way. In fiscal year 1996, ASCI funded eleven new code projects. Four of these code projects focus directly on increasing the ability of laboratories to predict the integrated performance and safety of weapons through computational means. The remaining code projects address materials aging and manufacturing issues and set the initial conditions for performance and safety codes.

On August 22, 1995, the Secretary of Energy announced the Defense Programs procurement of the world’s fastest computer, to be delivered to Sandia National Laboratories in December 1996. This computer will run at more than 1 trillion operations per second.[1] On February 20, 1996, the Department, through Los Alamos and Lawrence Livermore National Laboratories, announced a competition for a follow-on system that will explore a different approach to high-performance computing. It will run at 3 trillion operations per second. To support collaboration among ASCI researchers at the three laboratories, the first secure high-speed data network linking the Los Alamos, Sandia, and Lawrence Livermore National Laboratories was established on October 20, 1995.

The ASCI program will continue to advance during fiscal year 1997. Significant efforts will be required to continue the development of codes started in fiscal year 1996. The development plans for these codes include the delivery of initial capabilities needed to certify planned near-term stockpile modifications (for example, W76 recertification, W88 pit rebuild, and B61 Mod 11). The development of these codes also requires a significant interface with Defense Programs experimental facilities for verification and validation. Fiscal year 1997 will also see the delivery of two 100-billion-operations-per-second development systems at Los Alamos and Lawrence Livermore for the 3-trillion-operations-per-second computer.

Conclusion

In August 1995, the United States took a significant step to reduce the nuclear danger. The decision to pursue a “zero yield” Comprehensive Test Ban Treaty will allow greater control over the proliferation of nuclear weapons and will limit the growth (if not reduce the size) of the nuclear arsenals. This step is only possible because of the Science-Based Stockpile Stewardship program, which provides an alternative means of ensuring the safety, performance, and reliability of the United States’ enduring stockpile. At the heart of the SBSS program is ASCI, which will create the high confidence simulation capabilities needed to integrate fundamental science, experiments, and archival data into the management of the actual weapons in the stockpile. ASCI will also serve to drive the development of simulation as a national resource by working closely with the computer industry and universities.

1. For this document, operations per second, or OPS, will mean the scientific “floating point operations per second.”

1. Introduction

The Accelerated Strategic Computing Initiative (ASCI) is a focused and balanced program that will extend the Department of Energy's computational and simulation resources to create virtual testing and prototyping capabilities for nuclear weapons. By 2010, the ASCI program will—

- Develop high-performance, full-system, full-physics predictive codes to support weapon performance assessments, renewal process analyses, accident analyses, and certification
- Stimulate the U.S. computer manufacturing industry to create more powerful high-end supercomputing capability required by these applications
- Create a computational infrastructure and operating environment that makes these capabilities accessible and usable

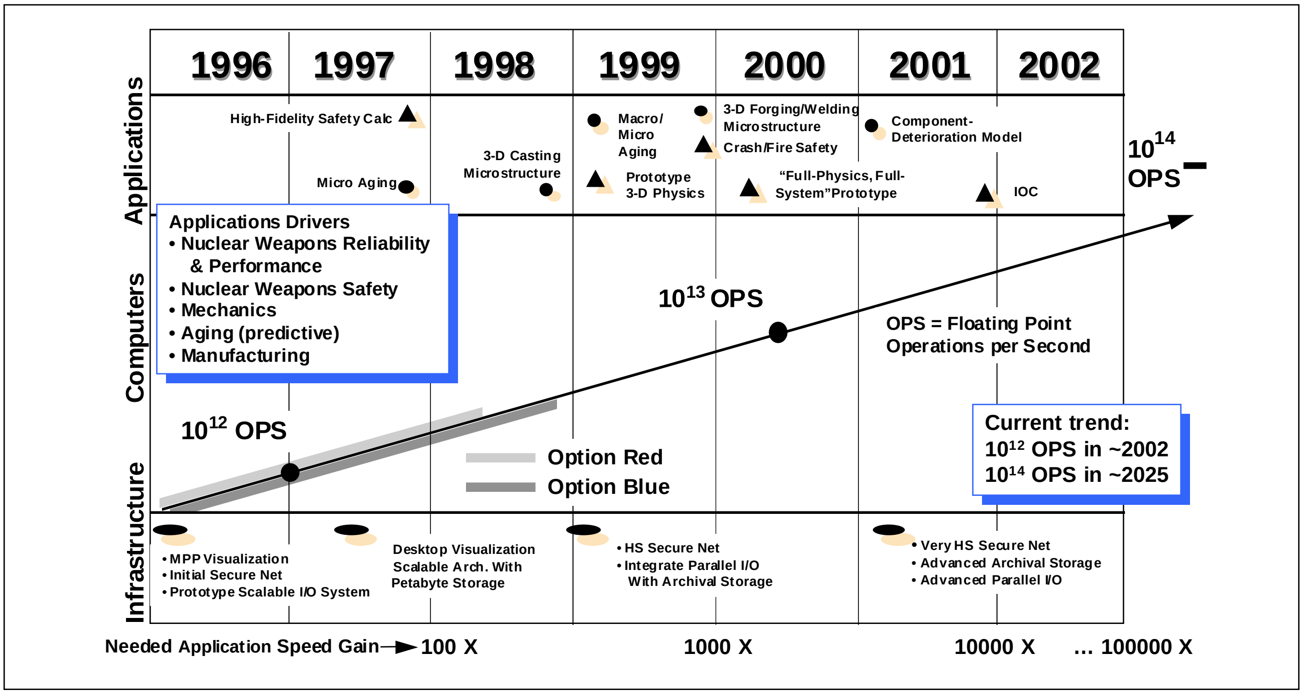

Figure 4. *A balanced program in applications, computers, and infrastructure will achieve needed performance gains.*

The United States' decision to pursue a true "zero yield" Comprehensive Test Ban Treaty, which President Clinton announced on August 11, 1995, marks the beginning of a new era for the nuclear weapon program. Over the last 50 years, the weapon program relied heavily on nuclear testing to ensure the reliability and safety of the nuclear stockpile. This approach included—

- Multiple tests per system
- Many prototype and exploratory designs
- Large numbers of systems
- A large standing complex
- Frequent modernizations

The President's announcement is the culmination of 5 years of work to move away from the former approach. In the future, the Department of Energy must maintain confidence in the performance, safety, and reliability of the enduring stockpile in an era in which the United States plans—

- No underground nuclear testing
- No new weapon designs
- Fewer facilities
- Reduced funds
- Fewer people

This environment calls for a new approach that relies heavily on the use of predictive simulations. Fortunately, dramatic advances in computer hardware and software technologies are now making virtual testing and virtual prototyping viable and practical alternatives to the traditional test-based methods of stockpile stewardship and management. ASCI will usher in the applications and technologies needed to promptly shift from nuclear test-based methods of assurance to computational-based methods.

Virtual testing and prototyping rely primarily on advanced numerical simulation capabilities to analyze, manufacture, and certify the performance of a nuclear weapon. Rudimentary versions of virtual testing and prototyping exist today. For example, it is now possible for an experienced weapon designer to perform two-dimensional simulations of the response of a weapon in certain abnormal environments using detailed

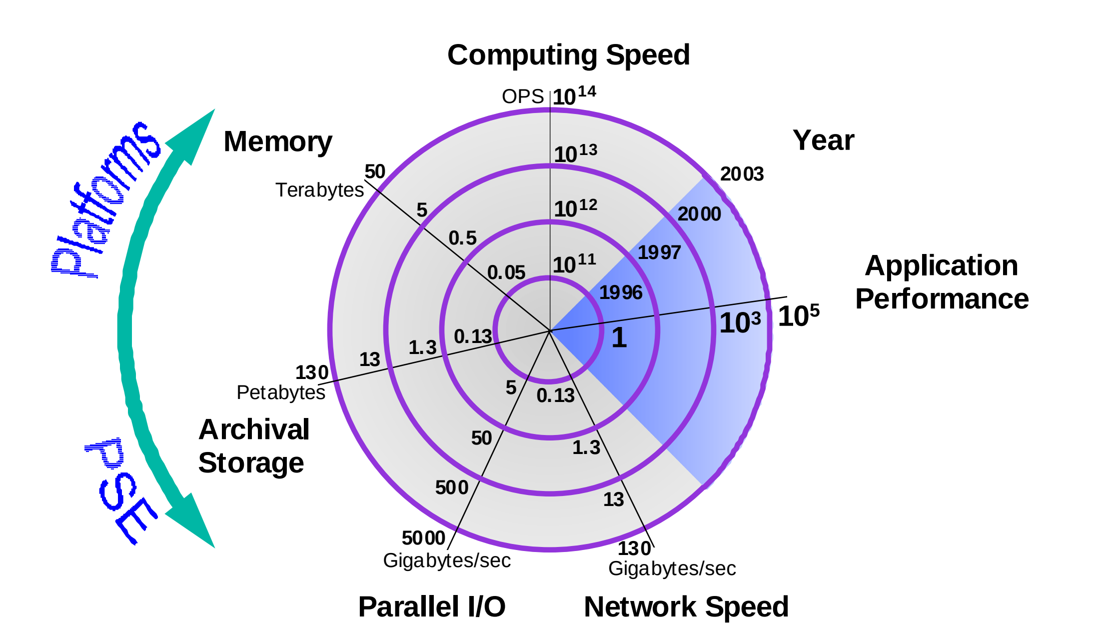

Figure 5. *ASCI success requires balanced growth of simulation components driven by the need for improved application performance.*

geometries and relatively high resolution based on empirical models.

However, to meet the needs of stockpile stewardship and management by the year 2010, applications must achieve progressively higher resolution, three-dimensional, full-physics, and full-system capabilities. The term “full-system” refers to the entire nuclear weapon system, including both nuclear and nonnuclear components. The term “full physics” implies a reliance on a suite of capabilities, with a firm theoretical and mathematical basis, that accurately captures the phenomena of interest, as opposed to capabilities based on empiricism. It does not necessarily imply a complete solution of the most fundamental physics in the actual applications code. In this context, “physics” is a generic term for the fundamental scientific and engineering modeling that must be included in the application codes; it includes physics, chemistry, materials science, and engineering sciences.

This level of simulation requires high-performance computing far beyond what is currently available. Current simulation applications are limited by today’s computing capabilities and rely heavily on underground testing. Current capabilities are not adequate for a detailed predictive analysis of unanticipated changes in a weapon design. In addition, these applications were not designed to operate efficiently on the high-end scalable architectures that will be required for virtual testing. These applications must be rewritten to incorporate refined physics algorithms, three-dimensional capability, and efficiency on advanced platforms.

2. Background

Test-Based Stockpile Stewardship

The United States has designed and maintained a stockpile of nuclear weapons for more than 50 years. Over that time, the United States Government, through its national laboratories and production facilities, developed methodologies to maintain confidence in the performance, safety, and reliability of nuclear weapons. These methodologies, both nuclear and nonnuclear, were generally test-centric. Scientists and engineers would attempt to apply the most complete physics understanding possible to designs or questions about the stockpile. Many times this would result in extensive mathematical predictions of a weapon's performance. As computer power increased, these predictions were incorporated into computer programs, which provided a higher degree of information to weapon designers. Because the physics understanding of the weapons was not complete, many empirically derived factors were incorporated into the computer codes to improve the codes. This led to a strong interdependence between the use of computers and testing.

In the early days of the weapon program, the national laboratories consistently purchased the highest performance computers in the world. These computers were needed to continue to improve the designs of nuclear weapons, making the weapons smaller and lighter, while improving safety and reliability. The computational power of these early computer systems was limited, and therefore codes continued to be one or two dimensional, with many empirical factors. That was sufficient in an era where extensive testing was conducted. The computer codes would predict the test results, and then the test results would be used to make specific calibrations to the codes for each weapon. In this situation, code limitations were obviated by the use of tests, which could serve as the final integrating factor.

Figure 6 shows the high-level steps in the detonation of a nuclear weapon. In the past and future, the scientists and engineers at the national laboratories would provide the best understanding of the physics for each step. This work was (and continues to be) supported by a wide array of aboveground test facilities and laboratory-scale experiments. The significant changes that affect the old methodology to provide confidence in the stockpile are the decision to pursue a zero-yield Comprehensive Test Ban Treaty and the loss of nonnuclear testing facilities. ASCI is intended to replace test-centric methodologies with computation-centric methodologies. This is not to say that the types of testing that are allowed under treaties and are affordable will not be conducted. In fact, ASCI anticipates an increase in this area of treaty-approved testing. The difference will be the use of such testing to validate and verify the physics models in the codes that in turn are

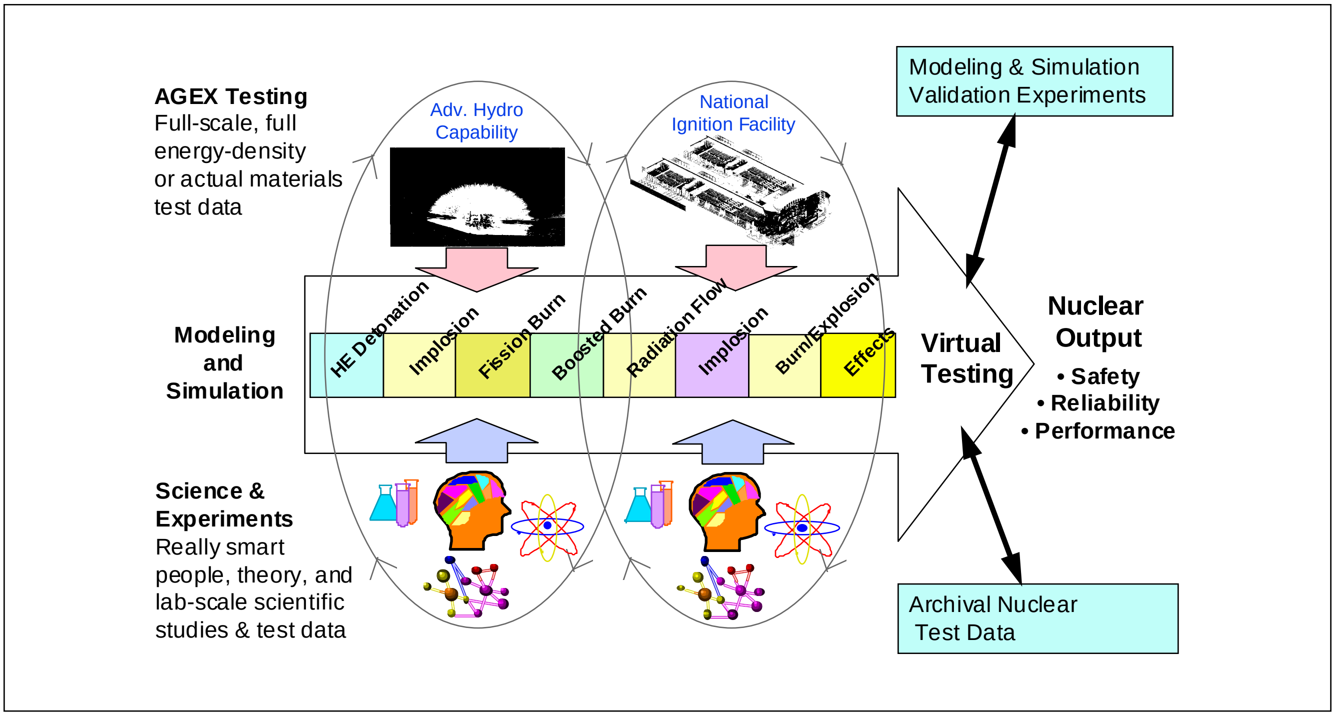

Figure 6. *The ASCI Program will enable the prompt shift from test-based methods to computational-based methods that are anchored to the stockpile through test data.*

used to predict the behavior of weapons in the enduring stockpile.

Not only will ASCI simulations, which are built on a large test base, compensate for the loss of underground testing, they also will mitigate the impacts of aging on stockpiled weapons. These weapons are now expected to be in service well beyond their original design lifetimes, making the systems susceptible to serious aging problems never before experienced. ASCI weapon models will be used to anticipate these failures, allowing for a balanced and planned approach to preventive maintenance and manufacturing rebuilds.

Teaming With the U.S. Computer Industry

The first supercomputers were developed for weapon applications in the 1960s as a partnership between the computer industry and the nuclear weapons laboratories. Throughout the 1970s and 1980s, this relationship continued. The Defense Programs laboratories were the early user sites and a primary customer for new state-of-the-art high-performance computers and computing simulation capability.

In the late 1980s and early 1990s this relationship changed. The Defense Programs laboratories have drastically reduced their partnerships with industry in the development of computer systems. As a result, the computing industry no longer views the Defense Programs laboratories as a primary customer for their most advanced computer designs. Moreover, the current test-based weapon simulation capability does not require the highest performance and largest memory machines. Consequently, the laboratories' latest computer acquisitions have not been the highest performance computers available.

The U.S. high-performance computer companies have responded to this decline by shifting their focus away from the scientific computing market and toward the business high-performance computer market, where the greatest demand is for cost-effective, midrange performance computers. While the industry is constantly improving the performance of its high-performance computers, the emphasis is on the price-to-performance ratio rather than on high performance in scientific applications. The result is that today the rate of improvement in high-performance computer speed is not keeping pace with Defense Program's new requirements.

Without a focused development program, the high-performance computers required for virtual nuclear testing, prototyping, and analysis of complex systems may never be developed. ASCI can promote the development of these peak high-performance computers by providing a visible and reliable market for these systems and by once again becoming a partner with industry in their development.

Teaming With Universities

Historically, universities have always had a close relationship with the Defense Programs national laboratories. In fact, Los Alamos and Lawrence Livermore have been operated for the Department of Energy by the University of California for many years. There have been times, however, when the missions of the universities and the Defense Programs laboratories have not been closely aligned. The Defense Programs laboratories' mission is focused on nuclear weapons, which requires very tight control of scientific information. Universities, on the other hand, generally encourage the free and open exchange of ideas and scientific knowledge. While this difference in the approach to handling information has sometimes led to tensions, ASCI and universities share a new common and critical interest. The success of ASCI depends on the ability to show that simulations can credibly be used as a proxy to replace testing as a means of ensuring stockpile confidence. Universities share this interest in proving that simulations can credibly reflect reality. Simulation has already proven valuable in exploring new scientific ideas. In fact, some are predicting that digital proxies for reality will create a new branch of science.

3. Program Strategies

The ASCI program has four main objectives in the areas of performance, safety, reliability, and renewal (see box). Each of these objectives requires computational capabilities that do not currently exist. Five strategies have been identified for meeting program objectives:

ASCI Objectives

- ***Performance:*** Create predictive simulations of nuclear weapon systems to analyze behavior and assess performance in an environment without nuclear testing.
- ***Safety:*** Predict with high certainty the behavior of full weapon systems in complex accident scenarios.
- ***Reliability:*** Achieve sufficient, validated predictive simulations to extend the lifetime of the stockpile, predict failure mechanisms, and reduce routine maintenance.
- ***Renewal:*** Use virtual prototyping and modeling to understand how new production processes and materials affect performance, safety, reliability, and aging issues. This understanding will help define the right configuration of production and testing facilities necessary for managing the stockpile throughout the next several decades.

- Create seamless management: one program—three laboratories
  - Operate ASCI as a single, three-laboratory program activity with seamless management and execution across the laboratories
  - Sponsor biannual Principal Investigator Meetings
  - Collaborate on development, and share hardware and software resources
  - Take maximum advantage of standard tools, common system structures, and code portability to enable interlaboratory collaboration
- Focus on advanced applications development
  - Focus on 3-D, full-physics, full-system applications
  - Focus on full-system, component, or scenario simulations
  - Validate simulations by rigorous correlation with constrained experiments and archival data
  - Accelerate code performance
- Focus on the high end of computing
  - Accelerate the development of scalable architectures
  - Develop partnerships with multiple computer companies to ensure appropriate technology and system development
- Create problem-solving environments
  - Provide applications development support

  - Ensure high-performance computing access
  - Develop tri-laboratory distributed computing environment
- Encourage strategic alliances and collaboration
  - Leverage other national initiatives
  - Collaborate with the best R&D programs of other DOE departmental offices, other agencies, universities, and industry
  - Attract the best researchers in the key disciplines for weapon applications

Together these strategies compose a balanced program that will create virtual testing and virtual prototyping capabilities.

Create Seamless Management: One Program—Three Laboratories

The problems that ASCI will solve for the Stockpile Stewardship and Management Program span the activities and responsibilities of the three Defense Programs laboratories. Cooperation is essential. The Defense Programs laboratories participate in the ASCI program as partners. This partnership strengthens the quality of ASCI deliverables and fosters synergy among the scientists and engineers at the laboratories.

Accordingly, the ASCI program has been developed with full participation from representatives of the Defense Programs laboratories. Project leaders from the three laboratories will implement the ASCI program with guidance from the Office of Strategic Computing and Simulation under the Assistant Secretary for Defense Programs.

Executing the ASCI program with the three Defense Programs laboratories acting in concert will require unprecedented levels of communication, coordination, and cooperation. Project teams will be formed to address specific milestones. Team members will be selected based on the skills the individuals bring to the effort, rather than their laboratory affiliations. The computational resources at each installation will be used by the researchers from all the laboratories. Scientists will test their specific applications on the various computing resources at each installation. As various technology paths are explored to achieve ASCI goals, the resulting products and prototypes will be available to all program members. To meet these goals, ASCI will define and develop a secure collaborative working environment among the laboratories.

Focus on Advanced Applications Development

ASCI applications are driven by two fundamental changes in the nuclear weapon landscape:

- The decision to negotiate a "true zero-yield" comprehensive test ban
- The requirement that weapon lifetimes extend well beyond their design lifetimes

The move from confidence based on nuclear testing to confidence based on predictive simulation forces a profound change in the questions the weapon designers ask of codes. In the past, they have used computations to address the question, "What should we test down-hole in Nevada?" Today, they must use the codes to answer the question, "Will this modified device perform as expected?" They no longer have recourse to actual underground nuclear testing.

Changes to stockpiled weapons arise from the aging process, design alterations, and changes in the manufacturing processes of replacement parts. When assessing the effect of these changes on weapon systems, weapon designers will use ASCI codes to develop new nonnuclear, aboveground experiments and integrate historical test data. To meet this challenge, ASCI codes require advances in the basic physics knowledge on which the codes rest as well as code technologies used in their implementation.

The Stockpile Stewardship and Management Program requires that Defense Programs extend the lifetime of existing systems into the indefinite future. This is complicated by three factors. First, the loss of key manufacturing technologies requires that we replicate existing functionality with newly designed components. Second, the reduction in size of the manufacturing complex requires that new and environmentally friendly manufacturing technologies be developed to produce necessary components. Finally, reductions in the overall budget for the weapon program constrain testing of both components and systems, even when this continues to be possible under current environmental regulations.

Traditional engineering approaches (design→ prototype→ test and evaluate→ reprototype→ manufacture) are not adequate for this challenge. Instead, ASCI will rely on predictive modeling and simulation for product realization. Comprehensive life-cycle engineering driven by high-fidelity modeling and simulation must replace traditional trial-and-error engineering approaches.

The progressively higher performance ASCI applications codes will provide a virtual test capability that is the key to achieving the Science-Based Stockpile Stewardship objectives outlined for 2010. The applications codes will provide the weapon analyst a "virtual test" capability as a means of developing a fundamental understanding of weapon system performance. The "virtual test" capability will provide a large measure of the confidence in the weapon systems that was previously attained by testing. Calculations by themselves cannot provide the data or empirical confidence in integrated system performance produced by nuclear testing. Weapon designers will use simulations (created with the new application codes and high-fidelity algorithms and physics models) in conjunction with archived test data and nonnuclear experiments to form the basis for weapon assessments and certification decisions.

It is a formidable challenge to replace the current empirical factors with predictive physical models. This challenge produces large, complex applications. These applications drive the scale of computing machinery and the infrastructure. Increased capability of machinery and infrastructure alone are insufficient. Much of the increased computational capability for ASCI must come from improvements in the applications codes themselves.

The use of improved physical models with the requisite fidelity requires applications that are much more complex than current codes allow. In addition, computer speed and memory must increase dramatically. In turn, a corresponding increase in the scale of the supporting infrastructure is required.

ASCI has three main thrusts for advanced applications development:

- Produce full-system, full-physics, 3-D simulations
- Rigorously validate simulations with experiments and archival data
- Accelerate code performance associated with ongoing stockpile investigations and problems.

Produce Full-System, Full-Physics, 3-D Simulations

Applications in weapon physics, engineering, materials, and manufacturing science will be the main focus of the simulation code development efforts. Together, they will provide the ability to predict the performance of full nuclear weapon systems in design analyses and in analyses of aging processes or complex accident scenarios. Performance simulations will cover the entire "stockpile-to-target" sequence to full thermonuclear yield and output.

These applications will also make it possible to design efficient and environmentally acceptable manufacturing processes capable of producing systems that meet nuclear certification and performance requirements.

Rigorously Validate Simulations With Experiments and Archival Data

Today's codes, constrained to run on current computers, rely extensively on approximations and empirically derived models. The new applications codes, incorporating models with a more fundamental physical basis and more accurate numerical algorithms, and running on much more powerful computing platforms, should produce results with much higher fidelity. Verification and validation of the component models and the overall codes are essential if weapon designers and analysts are to use them with confidence to address stockpile issues and certify weapon system performance without additional nuclear testing.

Verification and validation of applications codes are made more challenging because the codes will be used outside the range of previous test experience. Also, the codes will be applied to problems involving aging and component replacement processes and issues that were not confronted in the past. To provide the necessary confidence, we must demonstrate that—

- Underlying physical processes and phenomena are understood
- Numerical models and algorithms are faithful to the underlying physics
- Integrated predictions of the ASCI codes are reliable

Code verification begins with the code developers when they follow good software engineering practices and verify that the codes are written as intended. Ultimately, validation depends on direct involvement of the designers and analysts, who must accept the codes as useful tools. At a minimum, the validation process will include comparisons of code predictions against—

- Problems with analytic or well-established solutions
- Predictions of existing codes
- Data from aboveground experiments, component engineering tests, and full-systems tests
- Archived nuclear test data, where possible

Confidence is enhanced when a single code or simulation methodology can reproduce the results of a broad span of test data. Details of this process will be documented in the ASCI Verification and Validation (V&V) roadmap.

Accelerate Code Performance

The advanced computer systems developed by ASCI will incorporate leading-edge scalable (that is, parallel) hardware and software technologies and architectures, some based on entirely new concepts for high-performance computing. In general, the traditional weapon application development groups do not have the expertise for programming high-end scalable computers. As ASCI code teams are formed, scalable computer programming experts will be incorporated as full members of the teams.

The applications software being developed must not be hardware specific. Applications development will emphasize general numerical methods to operate on high-end scalable computer architectures. This will minimize rewriting of the new codes as architectural concepts change.

General numerical methods, however, are not usually as efficient as those optimized for specific architectures. Code acceleration also addresses efficiency of algorithms common to multiple codes, such as matrix solvers. These algorithms can be tuned to specific hardware architectures. The resulting set of algorithms can then be shared on a given platform by several applications codes. Applications code development teams will include experts in massively parallel processor programming and algorithm development, as well as experts in the physical processes being modeled. This allows the performance advantages of specific architectures to be studied and the results applied to ongoing code development projects and high-performance computing platforms.

In summary, the code acceleration project will—

- Assign programming experts as full members of the application development teams to produce more efficient algorithms and code structures
- Prepare model and simulation codes for use on ASCI platforms as soon as these platforms are available
- Study the comparative performance advantages of various computing architectures and programming methodologies
- Maximize the performance of specific applications on specific architectures consistent with good software engineering practices

Focus on the High End of Computing

ASCI applications will require a five-order-of-magnitude increase in computing performance above the sustained 50+ millions of operations per second currently provided by high-end commercial technology.[1] To support "full-physics," "full system" simulation, ultimately 100s of trillions of operations per second and beyond are required. The ASCI platform effort will bridge the gap between giga-scale and tera-scale computing through four strategies:

- Accelerate high-performance computing through multiple partnerships
- Develop software to take full advantage of hardware capabilities
- Maintain affordability by accelerating existing industry technology trends
- Provide balanced 100-trillions-of-operations-per-second computing platforms by 2003/4

Accelerate High-Performance Computing Through Multiple Partnerships With Computer Companies

This strategy pursues multiple research partnerships to develop alternative approaches to achieving 100 trillions of operations per second. Leading-edge development at the high end will be accelerated through different approaches at each of three times scales as indicated in Figure 7.

- Single computer company partnerships will provide a major increase in computing resources in 1 to 3 years to provide the computing resources necessary for further weapon applications development under ASCI.
- Consortia will seed industry R&D with ASCI-relevant projects to provide improved computing resources on a 5-year time scale.
- University research will develop the long-term improvements necessary to achieve the 2003/4 goal of 100 trillions of operations per second.

A cornerstone of ASCI is the long-term health of the U.S. high-performance computer industry that will provide the high-performance computers needed to support virtual testing and prototyping. Sustained development must build on economically viable corporate technology and business plans. ASCI will stimulate industry competition to adopt new technologies for high-performance computing.

The fundamental complexity of applications codes precludes a direct analysis to determine a favored computer architecture to support ASCI applications, and the diversity of physics modeling issues over many length and time scales requires the consideration of multiple architectures. Additionally, any

[1]For this document, operations per second (or OPS) will mean the scientific "floating point operations per second."

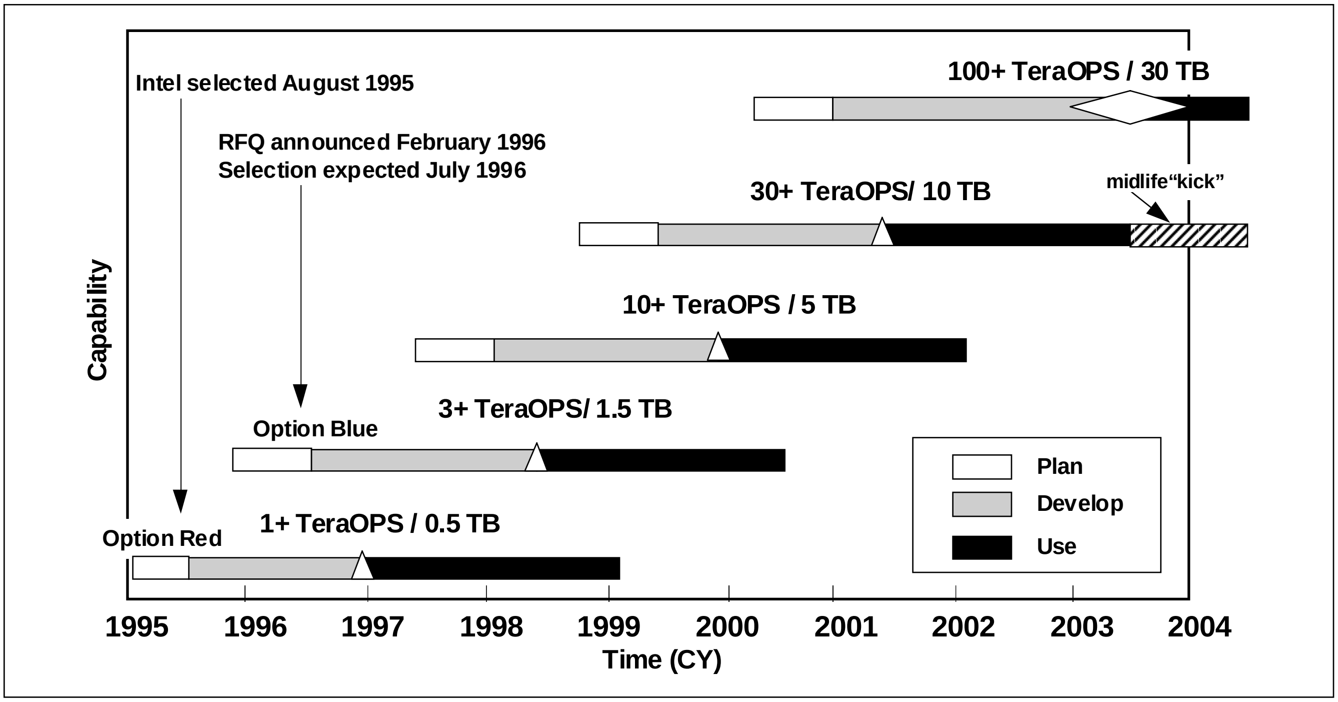

Figure 7. *ASCI computing systems roadmap—working with industry to reach unprecedented computer performance.*

single development program involves intrinsic risk. To reduce risk and to provide the breadth of computing capability required by the physics, ASCI will follow several promising technology approaches simultaneously. This will also allow competing technology and business aggregations room to develop and prosper.

Develop Software To Take Full Advantage of Hardware Capabilities

Of the 100,000-fold increase in computing performance in one decade that is expected by ASCI, a tenfold increase will come from improving software to take full advantage of parallel processing capabilities. This will be achieved in close cooperation with computer companies, consortia, and university researchers, with the Problem-Solving Environments strategy effort, to deliver software that will scale appropriately as hardware performance—speed, input/output (I/O), memory, number of processors, internal communications, and so forth—improves.

Collaborative efforts between the laboratories and computer companies will develop new software, data storage, and communications technologies to support high-end applications. Although several computer companies may participate in these partnerships, each partnership will specify unification of software environments to ensure portability of major ASCI codes between platforms.

Maintain Affordability by Accelerating Existing Technology Trends

The centerpiece of the ASCI platform strategy is to aggregate "commodity-off-the-shelf" technology in scalable architectures to develop high-performance computing platforms at affordable costs. The ASCI platform strategy will accelerate computer company business plans to create scalable high-performance systems from high-volume computer platforms. High-performance computing has become a niche market and no longer dominates development of computing technology. Instead, the enormous research and development effort of the computer industry is focused on this commodity market. It has led, for example, to a doubling of capability in microprocessor performance and dynamic random access memory (DRAM) density every 18 to 24 months for the past two decades at an almost constant price. The focus of the acceleration strategy, therefore, must be on the development of the overall architecture rather than on these individual "commodity" building blocks where the program can have little impact. At the same time, careful selection of technologies is necessary because the enormous scale of aggregation will multiply small unit cost differences into very large differences in system costs.

Provide Balanced 100 Trillions-of-Operations-per-Second Computing Platforms by the Target Date of 2003/04

The requirement for predictive simulations to assess the safety, reliability, and performance of the stockpile in support of stockpile life extension has led to the requirements to achieve platform performance objectives of 10 trillions of operations per second by the year 2000 and 100 trillions of operations per second by the target date of 2003/04. To support the day-to-day needs of weapon designers working in support of the Stockpile Life Extension Program, this requirement extends beyond just a burst-speed computation requirement. It requires a broad balance in both platform capabilities and the integrated problem-solving environments. This imposes requirements for appropriate memory, memory bandwidth and latency, disk capacity, parallel I/O bandwidth, and external networking.

Create Problem-Solving Environments

ASCI is an application-driven program. Its success depends on the rapid development of a new generation of weapon simulation codes and their application to the challenges of stockpile stewardship and management. The key players in these activities are the application developers and the weapon scientists. The fundamental goal of ASCI's Problem-Solving Environments strategy is to give these two groups the tools they need to do their job. This means—

- Providing the environment necessary for simulation codes that are adapted for the efficient use of very-large-scale parallel computers
- Ensuring that the power of the application-platform combination can be readily applied by DOE weapon scientists to the challenges of stockpile stewardship and management

To meet the goals of stockpile stewardship and management, ASCI will develop tera-scale computers thousands of times more powerful than the fastest computers in the world today. Making effective use of this power (that is, delivering it to the desktop) to solve the challenges of stockpile stewardship and management will require the development of a balanced, tera-scale support infrastructure also far more capable than any available today. All of the key components of this infrastructure must be scaled together with the growth in platform capability for ASCI to meet its goals. These are critical components that support the ASCI code teams. The development of this balanced support infrastructure is a challenge beyond

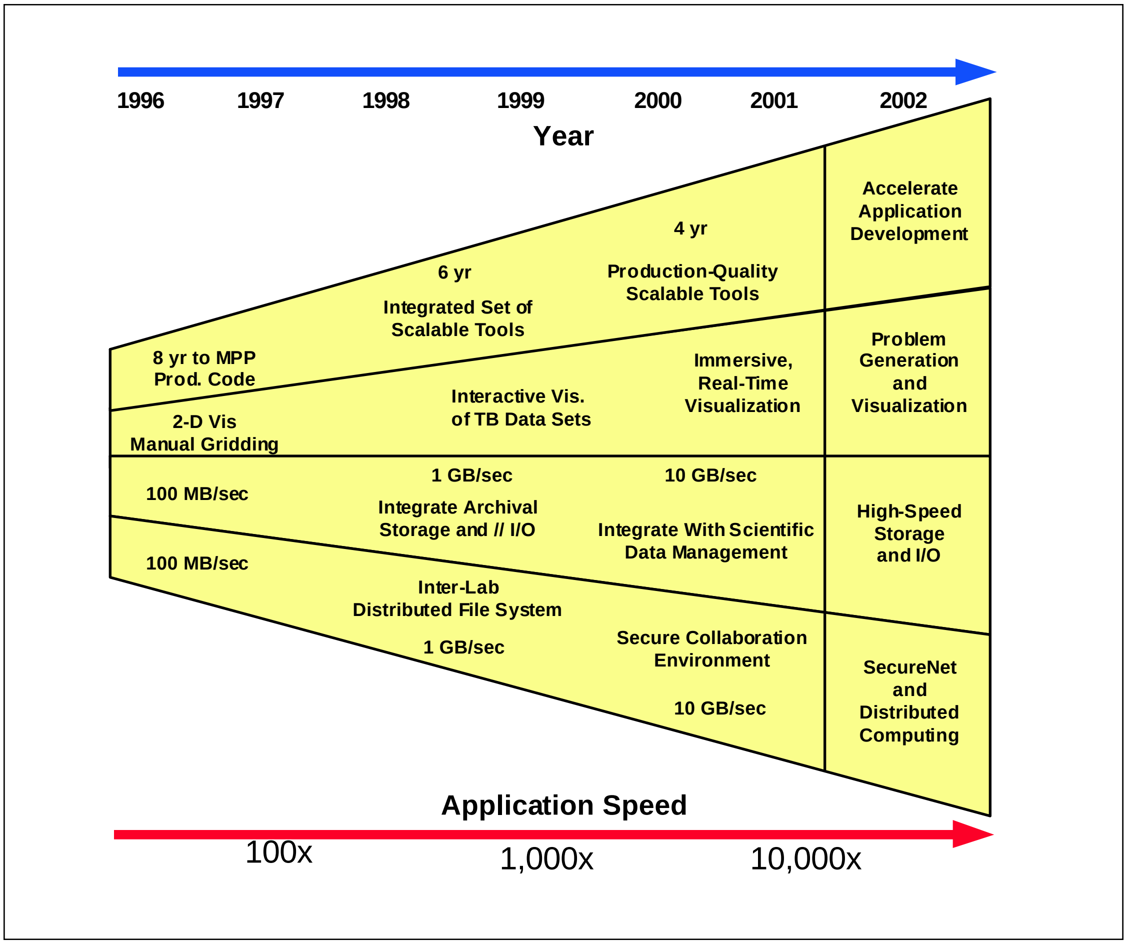

Figure 8. *Problem-Solving Environments Roadmap.*

the ability of ASCI, or any other single component of the high-performance computing community, to solve alone. ASCI will use existing and evolving products wherever possible. Where we must develop new products, we will seek to do so in collaboration with other members of the high-performance computing community.

The Problem-Solving Environments strategy has four main thrusts:

- Supporting code development for advanced weapon applications
- Supporting weapon application execution on ASCI computing platforms
- Ensuring that simulation results are accessible and understandable
- Providing a unified computing and information environment capable of supporting the Stockpile Stewardship and Management Program

Supporting Code Development for Advanced Weapon Applications

Developing and evolving ASCI's revolutionary scientific simulation codes to meet stockpile stewardship needs will be a challenging task. Today's large weapon codes typically take approximately 5 years to reach the point where they can be used with confidence by our weapon designers. Our experience with parallel code development indicates that, using the current set of tools, it will take 3 to 10 times as much effort to develop equivalently complex scalable codes. Hence, without the next-generation code-development tools, it could take up to 30 years before the new

ASCI weapon codes could be used with confidence by our weapon designers. Therefore, a major ASCI effort is to create an application development environment for the physicist/code developer that will reduce the parallel code development time to no more than that required for our existing codes.

The ASCI application development environment must provide the ability to rapidly develop complex applications that are efficient, scalable, portable, and maintainable. ASCI will work with academia and commercial computer companies to create the advanced development tools, methodologies, and standards needed to make this happen. This will require major advances in—

- Scalable software development tools for very large parallel codes
- Object-oriented parallel and distributed computing libraries and environments

Supporting Weapon Application Execution on ASCI Computing Platforms

To make effective use of the tera-scale code and platform capabilities being developed by ASCI, scientists will have to be able to store, access, and manipulate unprecedented amounts of data. Existing data storage, data management, system management, I/O, and networking technologies are inadequate.

Without orders-of-magnitude improvements in these areas, weapon designers will spend significant time and effort trying to work around data storage problems and I/O bottlenecks that will severely limit their ability to realize the enormous capabilities of the ASCI codes and platforms. This bottleneck situation has already been reached for the current generation of high-performance computers. ASCI's success requires a balanced environment where application performance and functionality scale with the increased capacity of the platform.

ASCI will eliminate these data bottlenecks by working with computer companies and other members of the high-performance computing community to achieve the needed orders-of-magnitude improvement in key infrastructure technologies. This will involve major efforts in the following areas:

- High-performance archival storage (including hardware)
- Scientific data management
- Very-high-speed parallel I/O and parallel file systems; and local area, high-speed interconnect fabrics

Ensuring That Simulation Results Are Accessible and Understandable

The capabilities developed in the ASCI program will be used by weapon designers to help make crucial judgments concerning the safety, reliability, and performance of the weapons in the U.S. enduring stockpile. Making good judgments will depend on their ability to interpret and understand the data available to them. Given the massive amounts of data involved, they will depend on graphically oriented data comprehension applications. These highly flexible applications must (1) allow the designers to directly examine all aspects of the simulation results, (2) provide powerful analytical capabilities with customizable and extensible human-oriented graphical interfaces, and (3) handle the massive amounts of data that will be generated by the ASCI code-platform combination. Such tools do not exist today. Their development will involve efforts in the following areas:

- High-performance visualization systems for data comprehension (including hardware)
- Graphical representation of large, multidimensional data sets

Providing a Unified Computing and Information Environment Capable of Supporting the Stockpile Stewardship and Management Program

The success of the Stockpile Stewardship and Management Program will require unprecedented levels of programmatic integration across the Defense Programs laboratories. Sharing information and computing resources among designers, code developers, and program managers will become the norm rather than the exception. Fundamental to resource sharing is secure, high-speed connectivity. A heterogeneous distributed computing environment will provide access to computing and information resources regardless of location while meeting security and need-to-know constraints. Tools capable of supporting large, time-critical collaborative efforts across the Defense Programs complex will also be required. The envisioned resource sharing requires—

- Wide area, high-speed classified networking connecting sites across the Defense Programs complex
- Seamless collaborative computing and information management environments for heterogeneous systems

Encourage Strategic Alliances and Collaborations

The shift to high-performance computing and science as the basis for confidence in the stockpile poses complex theoretical and practical problems in computer science and the physical sciences that are worthy of study by the best and most creative minds of the Nation. Engaging the efforts of individuals and groups at universities, other governments agencies, and industry through strategic alliances and collaborations will be critical to the success of ASCI. There will be three parts to this effort:

- The bulk of the Alliances strategy effort will be to form long-term strategic alliances with a small number of research universities and academic consortia to develop critical-mass efforts dedicated to long-term ASCI issues, such as high-confidence simulations.
- The Alliances strategy will also include smaller scale collaborations with individual investigators and research groups to work on more narrowly focused problems, such as turbulence.
- Task-oriented collaborations will be closely linked with specific ASCI deliverables and sponsored by laboratory research group managers to work on specific ASCI problems.

Collaborative work under this program must directly apply to weapon programs, yet be unclassified to the maximum extent possible. University work will be entirely unclassified. Examples of problems in high-performance computing include the following:

- Scalable architectures and I/O
- Algorithms
- Visualization
- Data management

The following are examples of problems in the physical sciences:

- Materials modeling, including the effects of aging
- Computational fluid dynamics, including turbulence modeling
- Modeling of manufacturing processes

Because of the unclassified nature of this work, it will be publishable in the open literature and may be presented at appropriate research conferences. This will serve several purposes:

- Subjecting algorithms, techniques, simulations concepts, and physical principles to broad expert review
- Ensuring the best and newest ideas are brought forth to address ASCI problems by circulating those problems among broad technical audiences
- Attracting and retaining the best minds in the program
- Enhancing the credibility, vitality, and validation of the program through broad and open interactions with the U.S. scientific community

4. Funding Plan

The ASCI funding plan is fully aligned with the ASCI strategies. The President's announcement of the decision to pursue a "zero yield" Comprehensive Test Ban Treaty and the requirement that weapon lifetimes extend well beyond their design lifetimes have placed additional demands on the ASCI program. These demands require a reestimation of the funding requirements for the program. This rebudgeting process is under way and will be included in the fiscal year 1998 budget request that is now being prepared by the Department of Energy.

The fiscal year 1997 Presidential funding request for ASCI is $121.6 million. Accordingly, the budget for the five ASCI strategies is as follows:

Applications	$54.9 million
Problem-Solving Environments	$23.5 million
Platforms	$33.7 million
Strategic Alliances	$ 6.1 million
One Program/Three Labs	$ 3.4 million

These funding numbers may change as ASCI's total budget is still subject to Congressional mandates in the upcoming fiscal year 1997 Senate appropriations bill. In such case, the new budget numbers of the individual strategies will be provided in the next version of this Program Plan.

5. Management Plan

Program Management Objectives

ASCI program management has two important objectives:

- **Provide leadership.** ASCI will develop computational capabilities that are critical to the Science-Based Stockpile Stewardship and Management Program. It must accomplish its mission on time and on budget. The ASCI management structure must enhance the work at the Defense Programs laboratories, keeping it focused on the ASCI mission. It must also ensure that resources are properly directed. The success of the program also depends on the computer industry and universities. The ASCI management structure must provide leadership to engage the computer industry and universities in the quest for simulation capabilities that can credibly replace testing as a decisionmaking tool.
- **Facilitate the interactions between laboratories.** ASCI is breaking new ground in the degree of collaboration required among the laboratories. The ASCI management structure must understand and break down artificial barriers that inhibit that collaboration.

Program Management Planning Process

ASCI program management will utilize the following planning process:

- **The ASCI Program Plan.** The Program Plan provides the overall direction and policy for ASCI. This plan serves as the strategic plan for the program and identifies the key issues and work areas for ASCI.
- **The ASCI Simulations Development Roadmap.** This roadmap is being developed to understand the basic needs of the codes that will be developed under ASCI. These needs include the basic science and models to support the codes, code features (for example, large problem sizes, adaptive meshes, and so forth) and data needed from experimental and archival data programs. This plan will also tie simulations development to other Defense Programs initiatives. Most importantly, the Simulations Development Roadmap will allow prioritization of code development efforts to be linked to actual stockpile concerns.
- **ASCI implementation plans.** These plans will be prepared annually and will describe the work planned at each labora-

tory to support the overall ASCI objectives. The implementation plans will be prepared by the strategy teams who will ensure that the work at each laboratory is closely coupled with the other labs. The implementation planning effort will begin in April and will conclude with a coordination meeting in July in anticipation of the beginning of the fiscal year on October 1.

Performance Measurement

ASCI managers realize that the successful development of modeling and simulation capabilities depends, in part, upon an effective and suitable performance-measurement program. Yet, measuring the overall impact of a collection of research and development activities has been frustrating because there is no widely recognized metric to measure their effectiveness. ASCI has adopted an approach that addresses this need by comparing actual output with planned output within a given year. This metric is referred to as the R&D Effectiveness Metric (EM). This method is not only consistent with the University of California approach to developing performance metrics, but also aligned with the Government Performance and Results Act of 1993.

Developing realistic, but challenging, project milestones is the first step toward producing a meaningful EM. ASCI's annual implementation planning process ensures that project milestones are both cost-effective and fully integrated with all other ongoing and planned activities. Project metrics within a particular ASCI program element, or strategy, are aggregated showing overall performance for that strategy. Progress over time is demonstrated through the use of indexed metrics.

Laboratory managers are responsible for both measuring and managing the performance of the projects within their purview. Performance is reported during ASCI project reviews at the semiannual Principal Investigators meetings.

Organization

ASCI's organization structure is designed to foster the focused collaborative effort required to achieve the program's objectives.

- Executive Committee. This body consists of two high-level representatives (a primary and alternate) from each laboratory and from Defense Programs headquarters. The Executive Committee sets overall policy for ASCI, develops programmatic budgets, and provides oversight for the execution of the program.
- Strategy Teams. These teams are responsible for the planning and execution (through the individual labs' management structure) of the implementation plans for each strategy (Applications, Platforms, Problem-Solving Environments, and Alliances). The strategy teams consist of two representatives (a primary and alternate) from each laboratory. The strategy teams are facilitated by representatives from the ASCI Headquarters (ASCI-HQ) team.
- ASCI Headquarters Team. This team consists of Defense Programs Federal employees supported by representatives from the Defense Programs laboratories. The ASCI-HQ team is responsible for ensuring that the program supports the overall Science-Based Stockpile Stewardship program. The ASCI-HQ team facilitates the program's interactions with other government agencies, the computer industry, and universities. Finally, the team sets programmatic requirements for the laboratories and reviews management and operating (M&O) contractor performance.

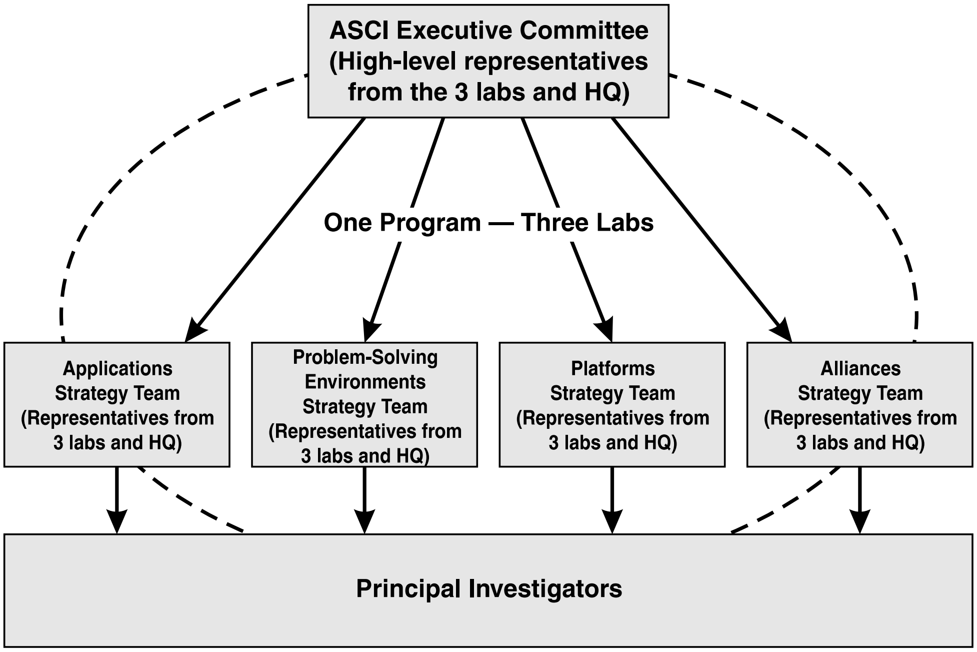

Figure 9. *ASCI organization.*

Program Collaboration Meetings

ASCI will regularly hold the following meetings as part of its leadership role and to facilitate collaboration among the three laboratories, industry, and universities:

- Principal Investigator meetings. These meetings will be held every 6 months and will provide a forum for ASCI Principal Investigators to meet and discuss progress on their research areas. These meetings will help foster collaborations by allowing principal investigators at one laboratory to present and discuss their work with their peers. The meetings (which will include participants from outside of the weapon labs) will also serve a peer-review function. These meetings will provide a semiannual technical review for the ASCI–HQ team.
- Executive Committee meetings. Every 2 weeks, the ASCI Executive Committee will meet via teleconference to discuss important issues. These meetings will ensure that relevant issues are identified, discussed, and resolved in a timely manner. The biweekly teleconferences will be supplemented with quarterly "face-to-face" meetings.

6. Relationship to the Stockpile Computing Program

The Defense Programs laboratories have an ongoing, highly regarded program to provide computational capability for current weapon analyses, engineering, and design efforts. This Stockpile Computing program provides the production simulation capabilities that support the enduring stockpile today and in the future, and it is the foundation for the long-range ASCI effort.

The programs have different goals, but they are complementary and interdependent. Stockpile Computing program goals include enabling, preserving, and advancing existing computational capabilities to meet stockpile support requirements. ASCI's goal is to produce simulation capabilities that meet the requirements of the designers and the stockpile and to transfer them to the Stockpile Computing program. ASCI will accelerate the transition of numerical simulation capability to a massively parallel environment by leveraging the skills and knowledge base of the Stockpile Computing program. The new high-performance applications in ASCI will

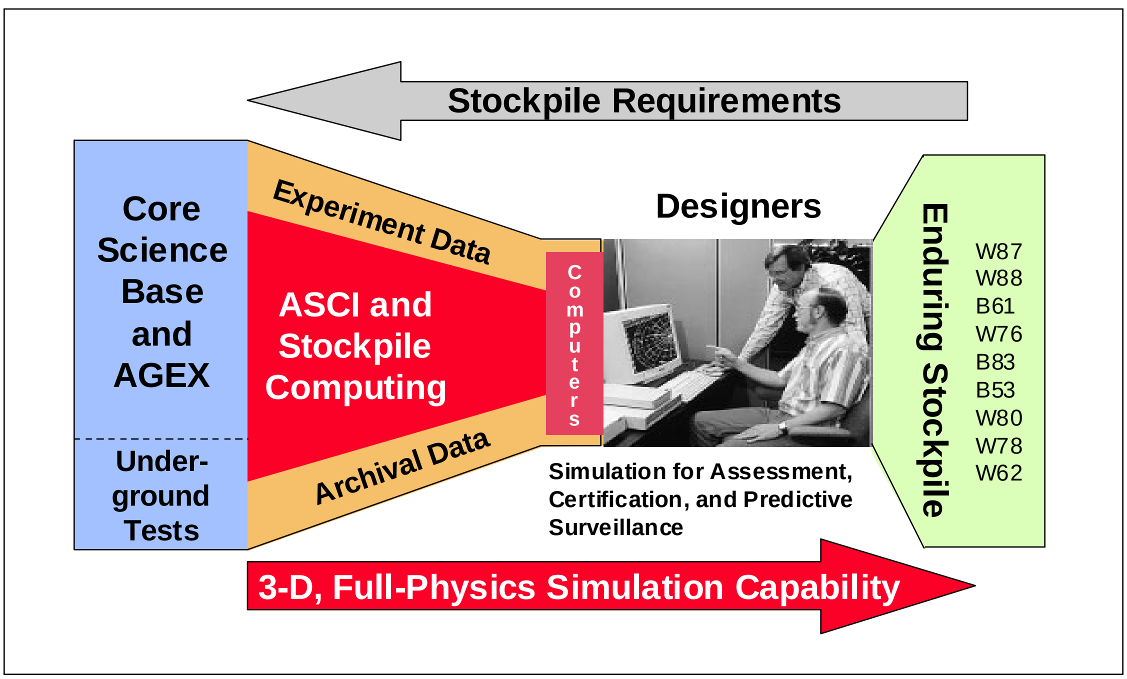

Figure 10. *ASCI provides the simulation capabilities to the Stockpile Computing Program to support weapon designers.*

be founded on the current, validated codes. Wherever possible, existing codes and algorithms will be directly applied to the new applications. Where modifications are required, the original developers will participate in the process to ensure that the changes preserve the intent and functionality of the existing models. Appropriate new physics models will be incorporated into the codes by experts who fully understand the limitations of the existing models. ASCI-developed applications will be maintained and applied by the Stockpile Computing program in the future.

ASCI expands the high-performance computing system technology and expertise available to the Stockpile Computing program. The ASCI research partnerships with computer companies will focus on research and development. Acquisition of production capability remains the responsibility of the Stockpile Computing program. The Stockpile Computing program will acquire the kind of high-performance technology that ASCI will create and prototype.

ASCI will provide the infrastructure required to support the availability and use of the ASCI high-end platforms for the development of application codes. The current computational infrastructure will not be replaced; rather, it will be extended and improved to meet ASCI's higher performance requirements. This new capability will tie the three laboratories together using a common architecture that will benefit the Stockpile Computing program.

ASCI is a program to extend and accelerate stockpile computing capabilities to meet new requirements. It will make use of existing stockpile capabilities and feed their results into the stockpile computing of the future. The success of ASCI depends upon a healthy Stockpile Computing program. The Stockpile Computing program must remain strong to apply the capabilities ASCI develops.

7. Relationship to the Stockpile Stewardship and Management Program

The shift from test-based confidence to computation-based confidence in the enduring stockpile has far-reaching consequences in Defense Programs. This shift requires [entails?] a new approach to the traditional methods used to ensure the performance, safety, and reliability of the nuclear weapons. Defense Programs has recently completed the *Stockpile Stewardship and Management Plan*, which describes the new approach to assessing and certifying weapons as they move from their current "full-scale," tested, and well-understood state to a state where they have aged, been remanufactured, or both. The ability to maintain confidence in the weapons as they move away from their current state will be provided through advanced simulation capabilities and improved scientific understanding.

ASCI must maintain a close relationship with both the Stockpile Life Extension Program (SLEP) and the Science Program. ASCI receives requirements from and provides capabilities to SLEP. ASCI in turn must

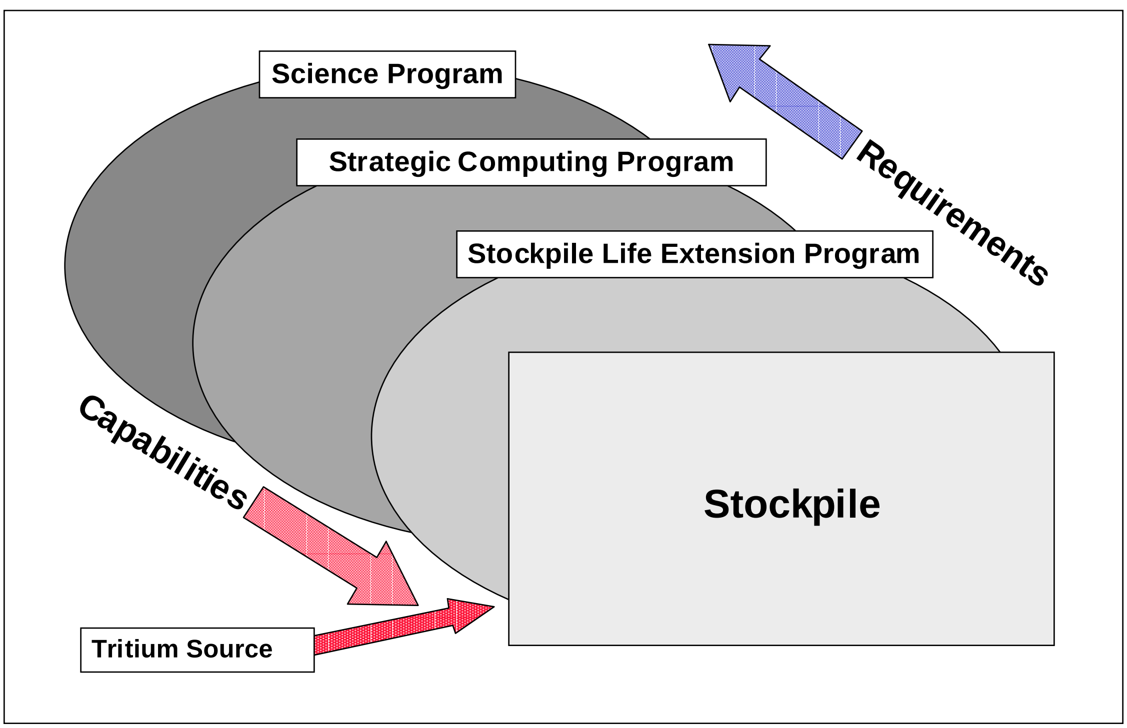

Figure 11. *ASCI is a critical enabler of the Stockpile Life Extension Program.*

look to the Science Program to provide additional science and experimental data to support the development of the simulation capability. ASCI will also provide important computational capabilities to the Science Program to support experiment design and to explore revalidation in a simulation environment.

The following is a short summary of the ASCI relationship with other elements of the Defense Programs Stockpile Stewardship and Management Plan.

Stockpile Life Extension Program

The Stockpile Life Extension Program is the operational arm of Defense Programs that conducts the surveillance, maintenance, refurbishment, and production activities for the stockpiled weapons. This program is also responsible for certifying and assessing the performance, safety, and reliability of aging weapons or weapons that are changed by refurbished parts that have been manufactured with new processes (for example, because of the inability to support the “old” manufacturing process due to cost or environmental concerns). ASCI supports SLEP by providing the advanced computational capabilities that allow these assessments and certifications to be made in the absence of testing.

ASCI will also provide important capabilities to address two major concerns with the stockpile. Today the United States’ nuclear stockpile is considered to be safe and reliable. In many cases, the weapons in the stockpile reached that safe and reliable state only after a period of time (shortly after the introduction of the weapon systems into the stockpile), when problems were identified and fixed. This is represented by the left side of the “bathtub” curve in Figure 12. Defense Programs currently anticipates that the weapons will remain safe and reliable for many years. The concern, however, is that the weapons will develop problems affecting their performance and safety because of old age (just like any electromechanical appliance, such as a photocopier or an automobile). Defense Programs is committed to identifying and correcting these problems long before they affect confidence in the stockpile (Figure 13). ASCI will provide vital computational capabilities to be used by the Enhanced Surveillance Program to predict when stockpile refurbishment is required.

The second major concern that SLEP must address is the introduction of new problems into the stockpile by refurbishment activities designed to rectify aging problems. ASCI will provide simulation capabilities that will predict how stockpile refurbishments will affect the performance, safety, and reliability of the weapons. These capabilities will be critical to keeping the number of new problems introduced by SLEP to an absolute minimum.

Stockpile Stewardship Program

The Stockpile Stewardship Program is the traditional source of science development, testing (including underground tests), and experiments needed to support the United States’ nuclear stockpile. In an era of test-based confidence, this program provided direct answers about the performance, safety, and reliability of the stockpile. The Stockpile Computing program was focused primarily at supporting the testing program. In an era without underground testing, with reduced aboveground tests, and with new manufacturing processes, the focus is shifting to computation-based confidence and the relationship between the Computing Program and Stockpile Stewardship.

In the future, Stockpile Stewardship will provide a critical foundation on which ASCI is built. Stockpile Stewardship will provide the understanding in science and the physics models needed to understand weapon perfor-

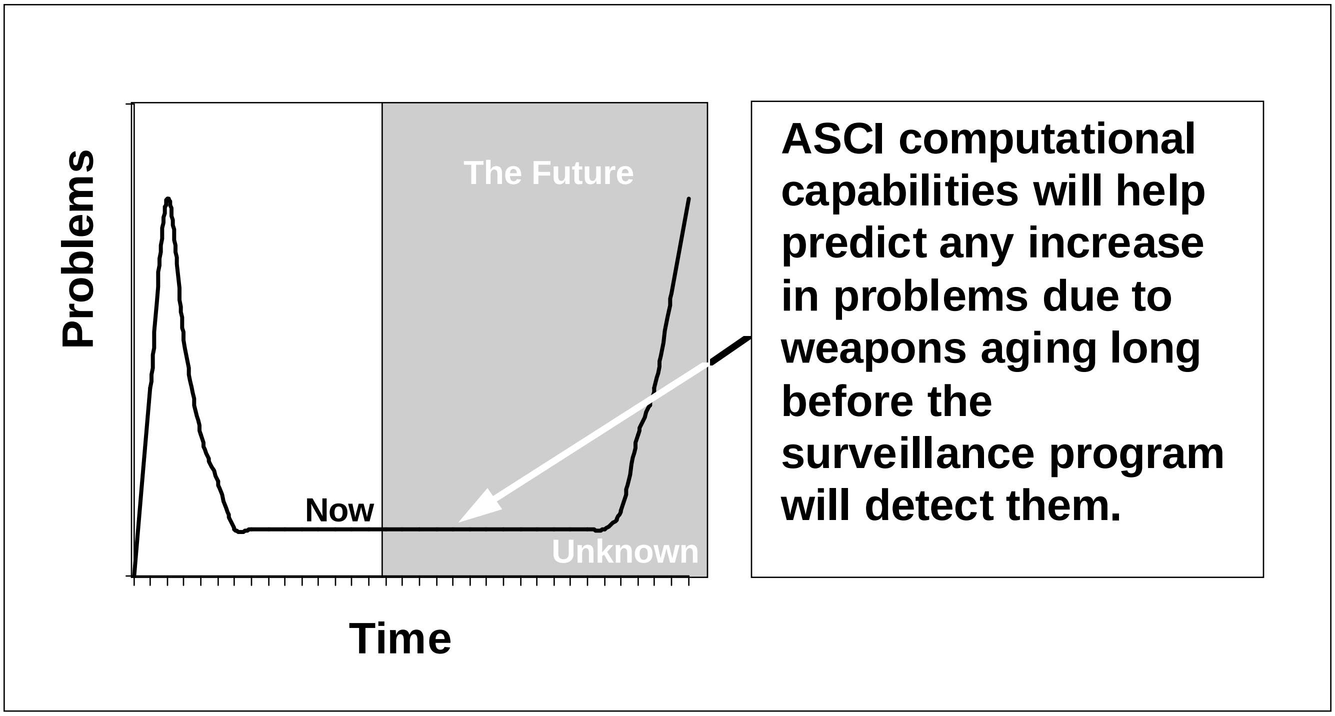

Figure 12. *ASCI supports SLEP by providing capabilities to predict aging problems.*

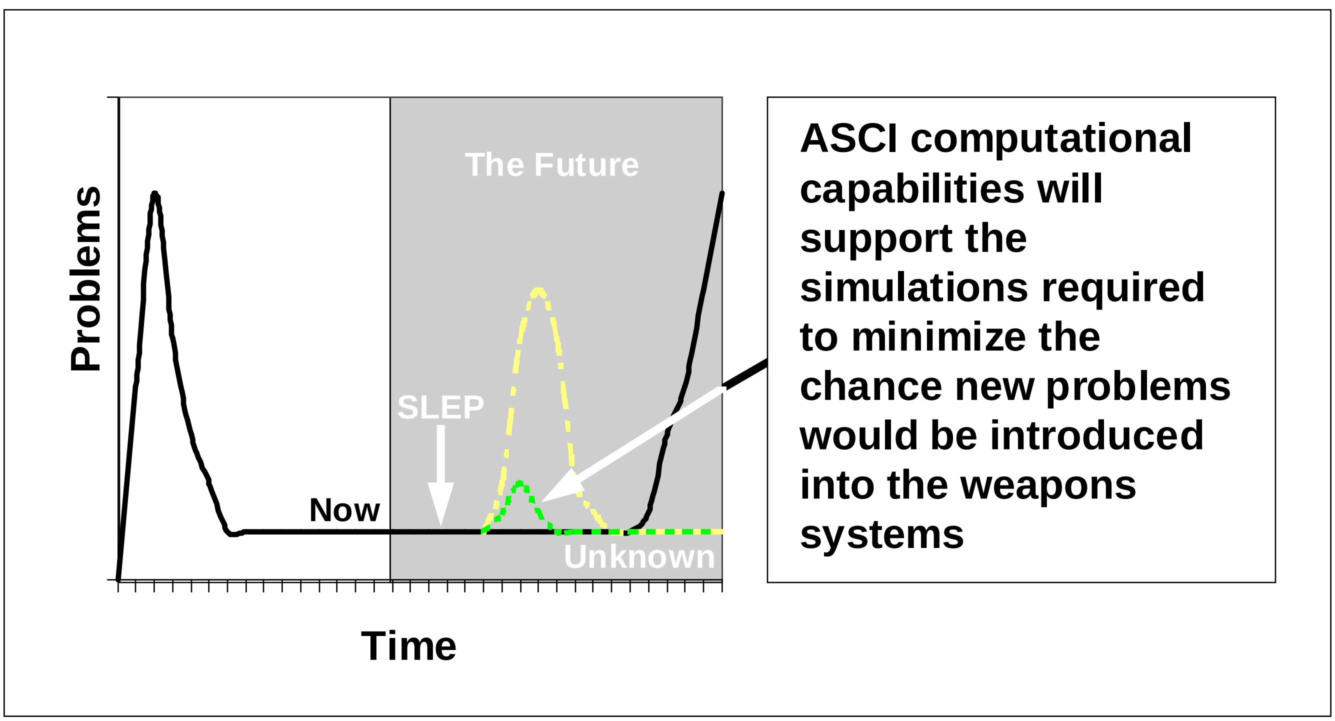

Figure 13. *ASCI will provide simulation capabilities needed to support Stockpile Life Extension refurbishment activities.*

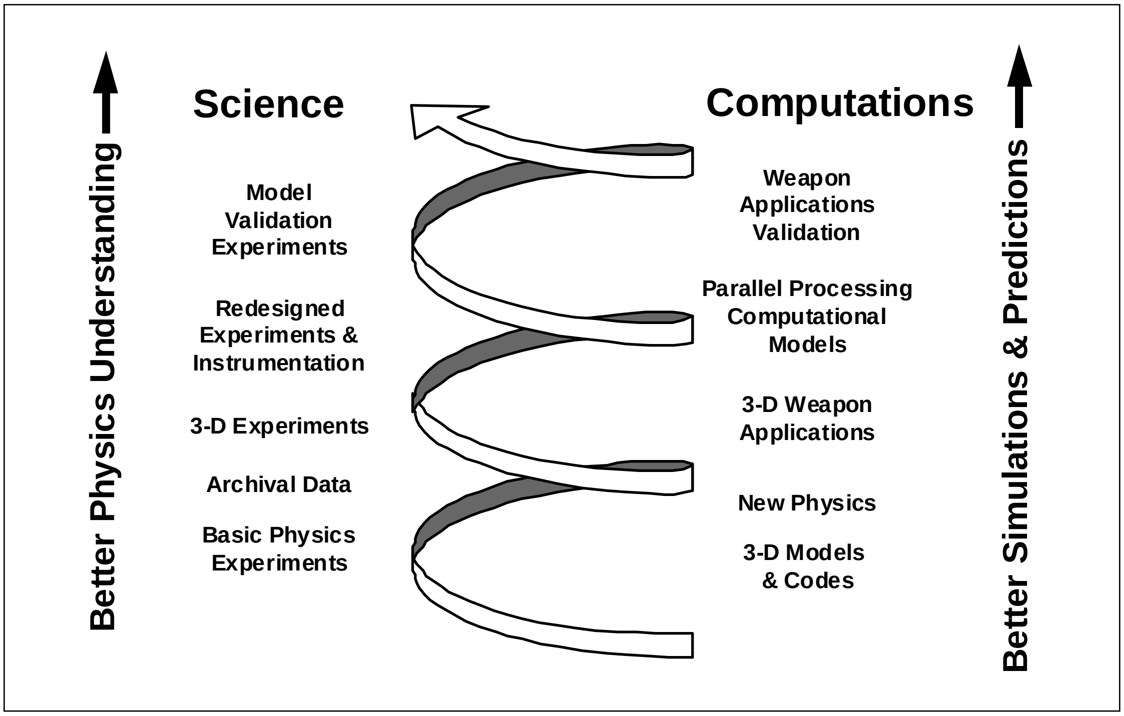

Figure 14. *ASCI is tightly coupled with the development of physics models in the Stockpile Stewardship Program.*

mance (Figure 14). Stockpile Stewardship will also provide the experimental facilities, like the National Ignition Facility and the Advanced Hydrodynamics Facility, which will be used to develop new physics models and provide validation data. Finally, the Stockpile Stewardship Program will maintain and continue to develop and understand the archival data from past underground tests, which provide a critical link back to full-scale weapon tests.

Acronyms & Abbreviations

3-D	three dimensional
ADaPT	Advanced Design and Production Technologies
AGEX	aboveground experiments
ASCI	Accelerated Strategic Computing Initiative
DOD	Department of Defense
DOE	Department of Energy
DP	Defense Programs
DRAM	Dynamic Random Access Memory
FY	fiscal year
GigaOPS	billions of floating point operations per second
HPC	high-performance computing
I/O	input/output
MegaOPS	millions of floating point operations per second
MPP	massively parallel processor
NIF	National Ignition Facility
OPS	Operations per second (for this document, OPS will mean the scientific “floating point operations per second”)
R&D	research and development
SBSS	Science-Based Stockpile Stewardship
SLEP	Stockpile Life Extension Program
SMP	shared memory processor
Tbyte	trillions of bytes
TeraOPS	trillions of floating point operations per second
V&V	verification and validation